

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

Requested Patent: WO0106362A1

Title:

METHOD AND SYSTEM FOR REGULATING BACKGROUND TASKS USING
PERFORMANCE MEASUREMENTS ;

Abstracted Patent: WO0106362 ;

Publication Date: 2001-01-25 ;

inventor(s): BOLOSKY WILLIAM J; DOUCEUR JOHN R ;

Applicant(s): MICROSOFT CORP (US) ;

Application Number: WO2000US18989 20000712 ;

Priority Number(s): US19990354970 19990716 ;

IPC Classification: G06F9/46 ;

Equivalents: AU6210500, EP1196848 (WO0106362) ;

ABSTRACT:

A method and system for regulating tasks of background processes so as to reduce interference with foreground processes. The progress rate of a background task (e.g., amount of work performed per unit time) is measured and evaluated against a target amount. If the progress rate appears degraded, the background task is suspended for a computed time interval so as to back off from its interference with a foreground process. Each time the progress rate appears degraded, the time interval is exponentially increased from its previous value up to a maximum, however if the performance appears normal, the time interval is reset to a minimum. Evaluation of the work is statistically based so as to eliminate variations in measurements, and automatic calibration of the target amount is provided, as is a mechanism for prioritizing multiple background tasks.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 January 2001 (25.01.2001)

PCT

(10) International Publication Number
WO 01/06362 A1

(51) International Patent Classification⁷: G06F 9/46

(21) International Application Number: PCT/US00/18989

(22) International Filing Date: 12 July 2000 (12.07.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/354,970 16 July 1999 (16.07.1999) US

(71) Applicant: MICROSOFT CORPORATION [US/US];
One Microsoft Way, Redmond, WA 98052 (US).

(72) Inventors: DOUCEUR, John, R.; 14705 N.E. 16th
Street, Bellevue, WA 98007 (US). BOLOSKEY, William,
J.; 24622 S.E. Mirrormont Drive, Issaquah, WA 98027
(US).

(74) Agent: MICHALIK, Albert, S.; Michalik & Wylie,
PLLC, 14645 Bel-Red Road, Suite 103, Bellevue, WA
98007 (US).

(81) Designated States (*national*): AE, AL, AM, AT, AU, AZ,
BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE,
ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP,
KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD,
MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD,
SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN,
YU, ZA, ZW.

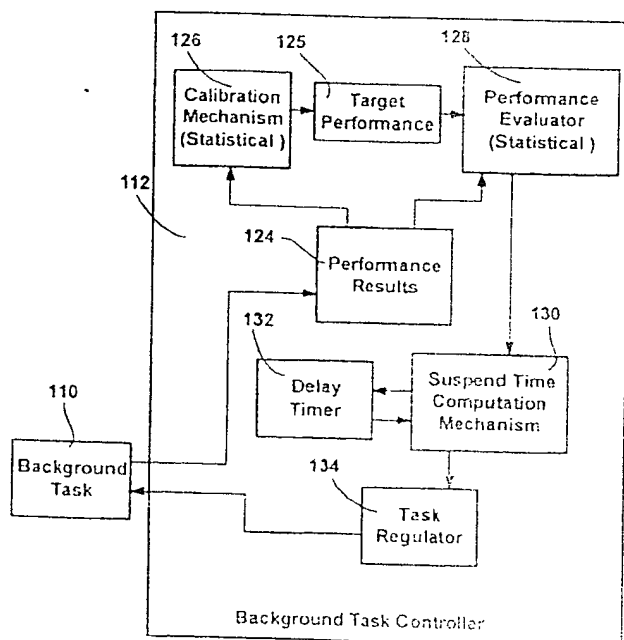
(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,
IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG,
CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

- With international search report.
- Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND SYSTEM FOR REGULATING BACKGROUND TASKS USING PERFORMANCE MEASUREMENTS



(57) Abstract: A method and system for regulating tasks of background processes so as to reduce interference with foreground processes. The progress rate of a background task (e.g., amount of work performed per unit time) is measured and evaluated against a target amount. If the progress rate appears degraded, the background task is suspended for a computed time interval so as to back off from its interference with a foreground process. Each time the progress rate appears degraded, the time interval is exponentially increased from its previous value up to a maximum, however if the performance appears normal, the time interval is reset to a minimum. Evaluation of the work is statistically based so as to eliminate variations in measurements, and automatic calibration of the target amount is provided, as is a mechanism for prioritizing multiple background tasks.

WO 01/06362 A1

METHOD AND SYSTEM FOR REGULATING BACKGROUND TASKS USING
PERFORMANCE MEASUREMENTS

TECHNICAL FIELD

5 The invention relates generally to computer systems,
and more particularly to improving the operating
performance of computer systems.

BACKGROUND OF THE INVENTION

10 Many computer systems execute background processes.
The tasks of these processes are often directed to
performing maintenance (housekeeping) functions, such as
file compression, disk defragmentation, file-system
content indexing and file archiving. While these tasks
15 are relatively important to complete at some time, there
is ordinarily no need to do so at any particular time.
Thus, they are often run in the background, operating
when no higher-priority foreground process, such as a
process of an application program interacting with a
20 user, is operating.

The traditional approach to operating a background
process is to assign the process a priority level that is
just above the system idle process, whereby CPU (central
processing unit) cycles are allotted to the background
25 process only when no normal priority process is ready to
use the CPU. For example, in the Unix programming
environment, the command "nice" initiates a process with
reduced scheduling priority, which may be set as low as
the system idle priority level.

30 Governing the execution of a process by priority
scheduling of CPU cycles assumes that the CPU is the
limiting resource. However, in many instances, the CPU
is not the limiting resource, but instead, process
performance may be limited by I/O (input/output) rate

and/or contention for another system resource, such as the disk, processor cache, memory capacity and/or memory bandwidth. Since background processes contend for these resources, background processes interfere with higher
5 priority foreground processes, and scheduling priority is insufficient as a mechanism for limiting this interference. The interference increases with background processes related to file-system maintenance activities, since such activities are resource intensive, but often
10 use few CPU cycles.

By way of example, consider a background process that wants to access the same I/O resource (e.g., a disk) as a foreground process, wherein both processes are regularly performing I/O operations to the disk resource
15 that take longer than the amount of time each process needs the CPU. When the foreground process is performing disk I/O operations, the background process is given some CPU cycles, during which time the background process desires access to the disk resource. As can be
20 appreciated, the background process has to wait for completion of the I/O by the foreground process.

When the foreground process completes its I/O operation, the background process is given access to the disk and the foreground process is given CPU cycles,
25 during which time the foreground process again wants access to the disk resource. However, because the foreground process has to wait for access to the disk resource, the foreground process has to idle, and at that moment, the CPU is essentially not being used by either
30 the foreground or background process. Accordingly, the CPU is not the limiting resource, but rather the disk is, i.e., contention for the disk resource by the background process causes the interference with the foreground process. Unless the background process is entirely

halted, however, the background process will receive CPU cycles when the foreground process is performing I/O, and thus will continue to interfere with the foreground process because the background process also will request
5 access to the disk resource during its allotted cycles.

A foreground application program's process that is of immediate consequence to a user is thus negatively impacted by (non-CPU) resource contention, unless the background process is canceled. However, simply
10 canceling the background process defeats the many advantages of having a background process run, e.g., to perform useful work during the many times when the foreground process is not contending with the background process for the CPU or other resource.

15

SUMMARY OF THE INVENTION

Briefly, the present invention provides a method and system for limiting the interference of a background process with a foreground process. A background process
20 is executed for a brief time slice. The actual performance of the background process is measured, and statistically analyzed to determine if the performance is degraded, operating normally, or if more information is needed. If the performance is degraded, the background
25 process is likely interfering with the foreground process, whereby the background process is then suspended (backed off) for longer and longer time intervals between executions until either some acceptable limit is reached or until the performance of the background process no
30 longer appears to be degraded, indicating that it is likely not interfering with another process. If normal performance is detected, the back-off time interval is reset to some predetermined minimum value. If normal performance is detected or more information is needed,

the task will again receive authorization to perform work. Because the actual performance is dynamically measured for each execution period, contention for a resource other than the CPU is actively detected, enabling the background process to appropriately and quickly back off from interfering with the foreground process via device contention. A critical task may have the suspension time dynamically adjusted based on its relative importance such that it will operate at a higher duty cycle.

The measured performance data may be used to automatically and statistically calibrate a target performance value for determining whether the measured performance is degraded. A mechanism for regulating multiple background tasks such that the tasks also do not interfere with one another is further provided. An exemplary application using the present invention to operate in the background to find and merge duplicate files is also described.

Other advantages will become apparent from the following detailed description when taken in conjunction with the drawings, in which:

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is a block diagram representing a computer system into which the present invention may be incorporated;

FIGS. 2A - 2B are block diagrams representing an example of a background process (groveler) suitable for use with the present invention that automatically locates duplicate files of a file system volume for merging;

FIG. 3 is block diagram representing various components of the exemplary groveler background process;

FIG. 4 is block diagram representing various components connected to a groveler worker object of FIG. 3;

FIG. 5 is a block diagram representing a general architecture for controlling the operation of a background task in accordance with an aspect of the present invention;

FIG. 6 is a block diagram generally representing components for regulating a background task in accordance with an aspect of the present invention; and

FIG. 7 is a flow diagram generally representing the steps taken to determine the operation of a background task in accordance with an aspect of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Exemplary Operating Environment

FIGURE 1 and the following discussion are intended to provide a brief general description of a suitable computing environment in which the invention may be implemented. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer. Generally, program modules include routines, programs, objects, components, data structures and the like that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers and the like. The invention may also be practiced in distributed computing environments where

tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

5 With reference to FIG. 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a conventional personal computer 20 or the like, including a processing unit 21, a system memory 22, and a system bus 23 that couples
10 various system components including the system memory to the processing unit 21. The system bus 23 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system
15 memory includes read-only memory (ROM) 24 and random access memory (RAM) 25. A basic input/output system 26 (BIOS), containing the basic routines that help to transfer information between elements within the personal computer 20, such as during start-up, is stored in ROM
20 24. The personal computer 20 may further include a hard disk drive 27 for reading from and writing to a hard disk, not shown, a magnetic disk drive 28 for reading from or writing to a removable magnetic disk 29, and an optical disk drive 30 for reading from or writing to a
25 removable optical disk 31 such as a CD-ROM or other optical media. The hard disk drive 27, magnetic disk drive 28, and optical disk drive 30 are connected to the system bus 23 by a hard disk drive interface 32, a magnetic disk drive interface 33, and an optical drive
30 interface 34, respectively. The drives and their associated computer-readable media provide non-volatile storage of computer readable instructions, data structures, program modules and other data for the personal computer 20. Although the exemplary environment

described herein employs a hard disk, a removable magnetic disk 29 and a removable optical disk 31, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, random access memories (RAMs), read-only memories (ROMs) and the like may also be used in the exemplary operating environment.

10 A number of program modules may be stored on the hard disk, magnetic disk 29, optical disk 31, ROM 24 or RAM 25, including an operating system 35 (preferably Windows NT), one or more application programs 36, other program modules 37 and program data 38. A user may enter
15 commands and information into the personal computer 20 through input devices such as a keyboard 40 and pointing device 42. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner or the like. These and other input devices are often
20 connected to the processing unit 21 through a serial port interface 46 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, game port or universal serial bus (USB). A monitor 47 or other type of display device is also connected to
25 the system bus 23 via an interface, such as a video adapter 48. In addition to the monitor 47, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

The personal computer 20 may operate in a networked
30 environment using logical connections to one or more remote computers, such as a remote computer 49. The remote computer 49 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all

of the elements described above relative to the personal computer 20, although only a memory storage device 50 has been illustrated in FIG. 1. The logical connections depicted in FIG. 1 include a local area network (LAN) 51 and a wide area network (WAN) 52. Such networking environments are commonplace in offices, enterprise-wide computer networks, Intranets and the Internet.

When used in a LAN networking environment, the personal computer 20 is connected to the local network 51 through a network interface or adapter 53. When used in a WAN networking environment, the personal computer 20 typically includes a modem 54 or other means for establishing communications over the wide area network 52, such as the Internet. The modem 54, which may be internal or external, is connected to the system bus 23 via the serial port interface 46. In a networked environment, program modules depicted relative to the personal computer 20, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

EXEMPLARY BACKGROUND PROCESS - THE SIS GROVELER

The present invention will be described below with respect to an exemplary background process that performs file I/O operations. In particular, one such background process, known as a single instance store (SIS) groveler, is directed to scanning a file system volume for duplicate files that can be combined into a single instance representation thereof. A single instance store facility and the SIS groveler are described in more detail in the copending United States Patent applications respectively entitled, "Single Instance Store for File

Systems" and *"Method and System for Automatically Merging Files Into a Single Instance Store,"* assigned to the assignee of the present invention, filed concurrently herewith, and hereby incorporated by reference herein in
5 their entireties. Notwithstanding, there is no intention to limit the present invention to any particular type of background process, nor to those that perform I/O operations, but rather, as will be understood, the present invention will have numerous uses and
10 applications throughout computing in general.

Turning now to FIGS. 2A - 2B, there is shown a general concept of a groveler 60 and a single instance store (SIS) facility and architecture, which may be implemented in the computer system 20 and which may
15 benefit from the present invention. As represented in FIG. 2A, in general, the groveler 60 runs as a background process to find files having duplicate data in a file system volume 62. Via a file system control named SIS_MERGE_FILES 64, the groveler 60 calls a Single
20 Instance Store (SIS) facility 66 to merge the duplicate files into a single instance of data with links thereto.

By way of example, as shown in FIGS. 2A and 2B, the groveler 60 finds that files 70, 72 (named Dir1\XYZ and Dir2\ABC) have duplicate data. When such duplicate files
25 70, 72 are located, the groveler 60 calls the SIS facility 66 via the SIS_MERGE_FILES control request 64. The call to the SIS facility 66 normally results in a single instance representation 74 of the original files 70, 72 with links 76, 78 (FIG. 2B) thereto, each link
30 corresponding to one of the original files, e.g., the user sees each link file as if it was the original file. The common store file 74 is maintained in a common store directory 80 of such files.

Each SIS link file 76, 78 is a user file that is managed by the SIS facility 66, while the common store 80 is preferably a file system directory that is not intended to be visible or accessible to users.

5 Ordinarily, the link files 76, 78 may be maintained on the same file system volume 62, as is the common store file 74 and the common store directory 80. To avoid a performance impact on other processes via pollution of the disk buffer cache, file reads by the groveler 60 are
10 non-cached. To avoid interfering with foreground processes via file locking conflicts, opportunistic locks are used by the groveler 60 when accessing a file, which temporarily suspend access to the file by another process until the groveler 60 can release it.

15 As generally represented in FIG. 3, to accomplish detection of duplicate files for merging purposes, the groveler 60 includes a central controller 82, which regulates the operation of one or more partition controllers 84_c - 84_e, one partition controller per file
20 system volume. In turn, the partition controllers 84_c - 84_e each have a groveler worker 86_c - 86_e associated therewith that when activated, individually attempt to locate duplicate files in their corresponding file system volume.

25 The central controller 82 synchronizes the operation of the partition controllers 84_c - 84_e across multiple volumes, for example, such that only one runs at a time depending on available system resources. In turn, when allowed to operate, each partition controller 84_c - 84_e
30 calls functions of its corresponding groveler worker 86_c - 86_e. As described below, the central controller 82 determines which of the volumes are allowed to be groveled, and in conjunction with each partition

controller 84, regulates how and when each groveler worker 86 is executed.

As represented in FIG. 4, each groveler worker 86 is a single process, and includes an open function 88, close
5 function 90, extract log function 92, scan volume function 94 and grovel function 96. In general, calling the open function 88 causes the groveler worker 86 to open (or create if needed) a database 100 of file information and a queue 102 of work items, also
10 conveniently stored as a database. The close function 90 takes care of any cleanup operations before the groveler worker is shut down.

The extract log function 92 uses a USN (Update Sequence Number) log 104 to add items (file identifiers)
15 to the work item queue 102. As is known, the USN log 104 dynamically records changes to a file system volume's files by storing change information indexed by a monotonically increasing sequence number, the USN. The extract log function 92 reads records from the USN log
20 104, each time starting from where it previously left off (as tracked by recording the USN), filters out those records that do not deal with new or modified files, and adds items to the work item queue 102 that correspond to new or modified files.

25 The extract log function 92 is called at a frequency that varies in an attempt to place a constant number of items in the work item queue 102 per call, and thus the frequency of calling the extract log function 92 is based on the amount of file system volume activity taking
30 place. For example, if disk activity is at a high rate, a large number of USN records will be extracted from the USN log 104, whereby the extract log function 92 will likely add a larger number of items to the work item queue 102 relative to times of slow disk activity. By

using the number of records extracted during the most recent extract log function call to determine the time duration before the next call, a high number of extracted records will cause a higher rate of calling extract log, while a lower number will cause a lower rate of calling. Over a period of time, the changes to the rate of calling roughly provide the desired number of items being placed in the work item queue 102 per call. Note that the rates may be adjusted gradually to smooth out any abrupt changes in disk activity.

The extract log function 92 records the last USN handled so that the extract log function 92 will begin at the correct location the next time it is called. As described above, the extract log function 92 returns a count of the USN entries extracted, from which the partition controller 84 calculates the next time to call the extract log function 92.

Calling the scan volume function 94 places work items (file identifiers corresponding to files in the volume) into the work item queue 102 via a depth first search of the file system volume 62. The scan volume function 94 is time limited, whereby when called, it places as many files as possible into the work item queue 102 within its allotted time, and resumes adding files from where it left off when called again. Note that the scan volume function 94 may be given some filtering capabilities, e.g., such that it will not add common store files to the work item queue 102, however at present the scan volume function 94 merely adds file identifiers as items to the work item queue 102 and any filtering is performed when work items are dequeued. The scan volume function 94 is called only when needed, e.g., when the work item queue 102 is first created, or if a problem occurs with the USN log 104 or the database 100,

since under these circumstances the extract log function 92 may not be provided with the proper file change information.

The grovel function 96 removes items from the work item queue 102, and processes each removed item to determine if it meets some criteria, e.g., whether the file corresponding to that work item has a duplicate file in the volume 62. To this end, the grovel function 96 computes a checksum (signature) from the file's data, and queries the file information database 100 via a database manager 106 (including a query engine) to see if one or more files in the volume have the same checksum and file size. At the same time, the database manager 106 updates the file information database 100 as needed with the file information, e.g., adds new records or changes existing records by storing or modifying the file size and signature indexed by the file ID in the database 100. If at least one matching file is found, the grovel function 96 performs a byte-by-byte comparison with the matching file set to determine if one of the files is an exact duplicate, and if exact, calls the SIS facility 66 via the SIS_MERGE_FILES control 64 to merge the files. In this manner, duplicate files may be automatically located and combined in a rapid manner by the groveler background process 60.

REGULATING BACKGROUND TASKS

Turning to an explanation of the present invention, FIG. 5 shows a general architecture wherein a background process 108 (e.g., a groveler) includes tasks 110₁ and 110₂, each of which is regulated by a background task controller 112 in a manner that substantially limits its interference with a foreground process 114. The architecture is made available as an application library

that provides a programming interface to monitor and control background tasks. Note that the actual unit of execution that is regulated is a task, of which there may be more than one per process. For example, in the SIS groveler 60, there are two tasks per disk partition, an "extraction" task (the extract log function 92) for moving extracted entries from the USN log 104 to the work queue 102, and a "grovel" task for the other groveler functions (e.g., the scan volume function 94 and the grovel function 96).

Although a single task may include multiple threads in order to maintain multiple contexts, in one implementation, only one thread is permitted to proceed at a time, in order to prevent threads from interfering with each other's work measurements. For example, in the groveler 60, there is one thread for dequeuing and processing files in a partition, and another thread for traversing the directory tree and enqueueing files, but both threads comprise the single grovel task. Task execution alternates between the threads (as determined by the process) in an attempt to maintain an approximately constant work queue depth.

In accordance with one aspect of the present invention, the background task controller 112 controls the activity of a process by separately throttling each individual task of that process. To regulate the operation of the background tasks of a process, each task (e.g., 110₁) waits for authorization to proceed from the background task controller 112. In general, the execution of each task is quantized into time slices, whereby each task is permitted to execute for a set period of real time. The authorization is given at a frequency that depends on a number of factors, including the performance of the background task, whether the task

is considered critical, and so on, as described below. A task may be given any CPU scheduling priority allowed by the system, e.g., normal priority, but may be throttled by the present invention by only allowing the task to
5 operate for limited time slices at a low frequency relative to how often a foreground process is able to request CPU cycles. A task may also be given a reduced CPU scheduling priority, whereby the CPU scheduling mechanism 116 will further control the cycles given to
10 the background task 110₁. As represented in FIG. 5, the background task 110₁ may be thus limited in how often it is given access to the CPU 21 and/or how often it obtains access to an I/O resource 118 (e.g., a disk via an I/O manager 120).

15 In accordance with another aspect of the present invention, to regulate a background task 110 by determining when it is permitted to execute, the background task controller 112 receives and analyzes performance results 124 (FIG. 6) from the background task
20 110, i.e., the progress or progress rate thereof, (e.g., amount of work performed per unit time). Based on these performance results with respect to a target performance amount 125, the background task controller 112 evaluates the work performed by the most recent background task,
25 and computes a time for suspending the background process to limit its interference with the foreground process 114. The background task controller 112 also uses the measured work performance 124 to statistically calibrate the target value for future comparisons. These functions
30 of the background task controller 112 are generally represented in FIG. 6 as components including a calibration mechanism 126, performance evaluator 128 and suspension time computation mechanism 130. After a delay timer 132 (e.g., a blocking mechanism) delays for the

calculated suspension time, a task regulator component
134 authorizes the background task 110 to again proceed,
whereby new performance results are obtained, analyzed
and used to determine a new suspension time for the next
5 authorization of the background task. This repeats as
long as the background process has a task to perform.

FIG. 7 represents the general steps taken to
calculate the suspension time, beginning at step 700
which represents the background task first being
10 authorized and executed. At step 702, the actual
progress, measured as an amount of work per unit time is
evaluated by analyzing the received performance results
124 against the target performance (target progress)
amount 125. As described in more detail below, the
15 progress (overall performance) is evaluated based on the
number of operations performed, the total time taken for
those operations, and the relative amount of work
performed by each operation. As also described below,
statistics are employed in the evaluation to deal with
20 fluctuations in performance measurements, whereby the
performance evaluator 128 returns either OK performance,
not OK performance, or unknown at this time, (i.e., needs
more information). In general, if performance (progress)
is good, then it is likely that the background task 110
25 was not contending with any foreground process, and thus
can continue its operations. Conversely, if performance
was poor, contention is suspected, and thus the
background task is suspended (backed off) by increasing
the time between its authorizations to proceed until
30 either performance improves or some upper limit is
reached.

As shown via step 704, if the performance was
acceptable (or better) with respect to the target
performance 125, step 704 branches to step 706 where the

suspension time is reset to some minimum. Note that when a throttled task's performance has returned within an acceptable target range, the throttling is removed by resetting the suspension time for the task to an initial,
5 minimum value. Step 708 provides an optional, (represented by the dashed box), working rate delay so that the background task 110 is not called too quickly, possibly reducing performance elsewhere in the system. Step 708 returns to step 700 to again authorize
10 performance of the task following this working rate delay.

If at step 704 the statistically-based evaluation mechanism 128 does not yet have enough information to gauge performance, step 704 branches around step 706 to
15 step 708. Since more information is needed by the evaluation mechanism, the task is permitted to proceed for another time slice, but the suspension time is left at its current value.

If at step 704 the performance was too low, step 704
20 branches to step 710 where a determination is made as to the importance of the background task. Important tasks may be analyzed against any appropriate criteria to determine when they should be treated as exceptions, whereby the task may become more aggressive by reducing
25 the suspension time and increasing its duty cycle. For example, the SIS groveler 60, which reclaims space by merging files, may be deemed critical as space begins to run low on a file system volume, such as below some threshold percentage. For regulated tasks of the
30 groveler 60, step 710 thus may include (at least occasionally) a check of the volume's free space. If a background task is deemed critical, the time between task operations may be adjusted (e.g., shortened) by step 712 as the relative importance of that background task

changes, regardless of its level of interference with any foreground processes. For example, as the limit on the free disk space approaches zero, the task suspension time approaches the working delay interval.

5 If at step 710 the task is not important enough to be an exception, step 714 delays the task for the current suspension time. If at step 710 the task is important enough to be an exception, step 712 adjusts the suspension time by reducing it before step 714 delays the
10 task for the suspension time. After the delay at step 714, the suspension time is increased at step 716.

 In one implementation, when poor performance is computed, the suspension time is doubled from its previous value, providing an exponential back off to
15 quickly reduce its interference with a resource-contending foreground process. The exponential increase, in combination with the reset when normal, makes the background task adjust to the time scale of the foreground execution patterns, providing relatively fast
20 background process resumption following short periods of foreground activity, but limiting the impact of the background task's execution during long periods of foreground activity.

 Step 718 adjusts the suspension time to ensure that
25 it is between some maximum and minimum value, e.g., if zero, doubling the time would have no effect, and if too long, the background task controller 112 would not be effective in detecting when the resource contention situation has improved. For example, the lower limit on
30 time between the slices of execution would be set to some minimum level (as described below), while the upper limit would bound the worst-case background task resumption time (as also described below). As described above, step 714 represents delaying for the suspension time, which

when complete, enables the background task controller 112 to loop back to step 700 to again authorize the background task to proceed.

Note that it is not necessary for all tasks of a background process to be throttled. For example, any task that primarily contends for the CPU may be sufficiently executed in the background by the existing CPU scheduling priority mechanism 116. Alternatively, if a task executes with a very low duty cycle, then the impact of that task on other processes may be very small, even if the task is I/O intensive during its active periods. For example, in the SIS groveler process 60, the extraction task 92 which extracts USN entries and queues them into the work item queue 102 is not throttled, since suspending it could cause missed file updates, whereby the groveler would have to rescan the entire volume. Instead, as described above, the groveler partition controller dynamically adjusts the execution frequency of the extract function 92 to attempt to maintain a fixed count of extracted USN entries per work queue update. This means that the duty cycle of the extraction task increases with disk activity, however even during periods of heavy disk activity this duty cycle has been measured at under two percent, which is deemed acceptable behavior.

As described above, when a task completes a slice of execution, the task reports back a set of performance data indicating the measured speed of its operations. However, the speed of a task may be highly variable, e.g., I/O times may vary widely even on an unloaded system, and clock granularity introduces errors in performance measurements. To overcome these fluctuations in performance measurements, the task controller 112

includes a statistical technique in the performance evaluator 128.

To this end, each task reports its performance data (progress) as a list of three components, i.e., triples, each parameter corresponding to a type of operation performed by the task. Each triple includes a count of the number of operations performed during the time slice, the total real time taken for the operations to complete and an operational parameter that represents the relative amount of work performed by each operation. The set of operations for each task is statically established prior to the regulated execution of the task, and operations may be common across tasks. In general, an operation may be any action or set of actions that consumes a sufficient amount of real time to be readily measurable. Operations that are CPU-intensive will exhibit performance that is sensitive to contention for processor cycles and processor cache, operations that are memory-intensive will exhibit performance that is sensitive to contention for memory capacity and memory bandwidth, and so forth. For example, in the SIS groveler 60, the performance of the process is primarily disk reads, whereby "disk reads" is the type of operation and the parameter is the amount of data read. Note that neither the background task controller 112 nor the process 108 needs to know the resources being used by an operation.

For each operation type, the operation time is computed as the total consumed time divided by the operation counts. The resulting operation time value is then compared against the target value (described below) using a paired-sample sign test, a known (Fruend, J.E., *Mathematical Statistics*, Fifth Edition, Prentice Hall, 1992) distribution-free, non-parametric statistical comparison test. The paired-sample sign test enables the

measurements from multiple operations to be combined into a single statistical test. More particularly, the paired-sample sign test compares each sample with a corresponding target value and counts the total number of samples that are greater than their target values. Each target value is a function of the operation type and the optional work parameter. The sample size increases linearly with each test, up to a set limit to bound the computational and storage requirements of the test, thereby taking advantage of more information as it becomes available.

The result of the comparisons may be expressed as an ordered pair (n, r) , where n is the sample set size and r is the count of samples that are greater than the target value. If r is greater than a threshold value that is a function of both n and a predefined (type-I) error probability α , then the performance is judged to be degraded, whereby the suspend time is doubled up to an upper limit as described above. If r is less than a threshold value that is a function of both n and a predefined (type-II) error probability β , then the performance is judged to be normal, and the suspend time reset. If r falls between the two threshold values, then the performance judgment is deferred, i.e., the performance evaluator returns its "do not know / need more information" result whereby the current suspend time is maintained at its previous value, as also described above.

A task is thus regulated based on whether its actual, measured performance is normal or degraded relative to a target value. A target value for a task may be manually determined, e.g., by empirically testing one or more samples when a machine is not loaded to determine a reasonable target performance for a task, and

then using that target in the comparison with the actual performance.

However, an improved way in which to obtain and maintain a target value is to perform an automatic calibration procedure using the same measured performance data used in the comparison by the performance evaluator. 128. To obtain a better representative target value 125, the actual values are conditionally sub-sampled, statistically abstracted and temporally filtered as described below.

For operations that perform a constant amount of work, the measure of the target performance is a quiescent operation time, which is the time to perform the operation on an unloaded resource. For operations that perform a variable amount of work, the time to perform each operation is assumed to correlate linearly with the quantity of work completed, such that the measure of the target performance is essentially a line that represents the quiescent operation time as a function of the work parameter.

Statistical abstraction produces a quantified measure of the target performance from the sampled performance measurements. This measurement is a linear equation that relates the operation time to a work parameter. The coefficients of the equation are obtained through least-squares linear regression, where the variable x is the work parameter, and the dependent variable y is the operation time. The six sufficient statistics for linear regression, $\sum x$, $\sum y$, $\sum xx$, $\sum yy$, $\sum xy$ and n are maintained persistently, whereby the performance target is not recalculated from scratch each time the process is executed.

The quiescent performance of an operation may change over time, e.g., the performance characteristics of file

systems change as successive iterations of writes cause files to fragment. Temporal filtering may be used to age out older performance data so that throttling decisions are not made based on stale data. Each time a sample is
 5 taken, the sufficient statistics are updated according to the following rules:

$$\begin{aligned}\Sigma_x &= \xi \cdot \Sigma_x + x \\ \Sigma_y &= \xi \cdot \Sigma_y + y \\ \Sigma_{xx} &= \xi \cdot \Sigma_{xx} + x^2 \\ 10 \quad \Sigma_{yy} &= \xi \cdot \Sigma_{yy} + y^2 \\ \Sigma_{xy} &= \xi \cdot \Sigma_{xy} + x \cdot y \\ n &= \xi \cdot n + 1\end{aligned}$$

wherein when ξ is strictly between zero and one (i.e., not
 15 including the endpoints of zero and one), these rules effect an exponential averaging of the regression state. The determination of ξ is described below.

Repeated application of the exponential-averaging update rule for n will cause it to asymptotically
 20 approach a constant value n' given by the following equation:

$$n' = 1 / (1 - \xi)$$

As a result, the effective steady-state sample size is
 25 n' . To initialize the regression, the calibration mechanism collects some number n'' of samples, weights them by the ratio n'/n'' , and calculates the sufficient statistics with $\xi = 1$. Since the background task controller 112 has no target performance values until
 30 after these n'' samples are collected, the background task controller 112 does not attempt to dynamically regulate the execution of the task during this initialization

phase. Instead, the background task controller 112 allows the task to execute at a special calibration rate that is both lower than the normal working rate, so as to reduce the performance impact of the task on the system, and higher than the suspension rate so as to complete the initialization phase within a reasonable period of time.

The exponential averaging of target performance is clocked by the rate at which samples are taken, rather than by the passage of time. Since measurements of degraded performance cause execution to back off, fewer such measurements will be made per unit time than will measurements of normal performance. Thus this mechanism will tend to capture values that reflect normal performance in preference to those that reflected degraded performance. However as described above, the execution rate of a task also increases as its work becomes more critical, even if its performance is degraded. Under these circumstances, the measurement data may be sub-sampled so as to consider only the number of samples that would have been taken if the maximum suspension rate had not been reduced.

Since the lack of target performance data is essentially what prevents the background task controller 112 from throttling a task during initial calibration, the completion of the initial data collection is sufficient to allow throttling to begin. However, there is no guarantee that the initial performance measurements were made during a time of relatively low system activity, so it is possible that subsequent performance that appears normal, by virtue of a comparison to the initially calibrated target, is actually degraded, whereby the background task controller 112 will fail to properly throttle the task. However, calibrating over a longer time increases the likelihood of encountering an

idle period so as to obtain accurate measurements of quiescent performance. A probationary period may thus follow the initialization phase, during which the execution rate is permitted to drop as low as the suspension rate when the performance is judged to be degraded, but is not permitted to exceed the calibration rate even when the performance is judged to be normal, since the normal judgment is suspect at this time.

As described above, several parameters control various aspects of background task execution. It is relatively straightforward to determine appropriate values for these parameters from desired properties of the background process. For example, each task may determine the length of its own execution time slice. By shortening the time slice, the task increases the promptness with which it can be suspended, however, the process needs to be long enough to efficiently amortize the overhead of the throttling mechanism. Tasks that perform operations with relatively long completion times will require longer time slices, e.g., in the SIS groveler 60, time slices for the groveling task are on the order of 200 milliseconds. The working rate defaults to unity, however this may be increased to limit the maximum possible performance interference of the background task on a foreground process.

The minimum suspend time is five seconds, but using a larger value delays the task longer before retrying, for more conservative behavior, while shorter delays permit a more rapid resumption of a task execution following a falsely triggered suspension. The maximum suspend time is five minutes, which keeps the impact below one-percent for the time-slice durations of most processes. As is understood from above, smaller values for the maximum suspend time reduce the worst case time,

while larger values reduce the impact on foreground processes.

The steady-state calibration sample size n' and the initial calibration sample size n'' are set to smooth out short-term variations in operation performance. The n' sample size value is used to calculate ξ (by inverting the above equation for n' , yielding $\xi = 1 - 1/n'$), and thus also determines the time constants of the calibration system. Performance measurements that indicate normal performance will be sampled at the working rate, whereby multiplying n' by the time slice duration yields the time constant for improved performance. The time constant for deteriorated performance is greater than this in proportion to the ratio of the suspension rate to the working rate. For the SIS groveler example, n' is set to 5000 and n'' is set to 1000, which for the 200-millisecond time slice and 5-minute maximum suspension time yield a performance increase time constant of 17 minutes and a performance decrease time constant of 17 days. Although the latter appears large, it is reasonable if no substantial and sustained decrease in disk performance is expected over a shorter period.

The calibration rate is generally set to a value between the suspension rate and the normal working rate. The higher the rate, the sooner the system will collect enough measurement samples to begin regulated execution. The lower the rate, the lower the performance impact on any foreground processes that may be running during target initialization. The default value is one time slice of execution per second, which for a 200-millisecond time slice limits the duty cycle of the background tasks to twenty percent. For an n'' of 1000,

this calibration rate will initialize the target performance in 3.3 minutes. The probation period may be set to virtually any duration, including zero, however a shorter period increases the danger of a negative impact on system performance due to an incorrectly calibrated target, whereas a longer period postpones the start of working at the maximal rate. The default value for this parameter has been chosen at twenty-four hours, so as to include any comparatively quiescent periods in a diurnally cyclical load.

The task duty cycle constant determines the period over which weighted fairness among tasks is maintained. By default, this value is set to five minutes to coincide with the default value of the maximum suspension time.

The two error probabilities (α and β) that govern the sensitivity of the statistical comparison test, are basically determined from three competing considerations. Increasing the value of α allows faster reaction to degraded performance, increasing β relative to α improves the stability of background execution, and decreasing β reduces the performance impact on foreground processes. A minimum number of samples is collected before degraded performance can be recognized by the paired-sample sign test. This minimum sample size increases with decreasing values of α . For the paired-sample sign test, the relationship between α and minimum sample size m is quantized negative logarithmic:

$$\lceil m = -\log_2 \alpha \rceil$$

Therefore, a minimum of m performance measurements is made before the performance can be judged to be degraded and the background task thereby suspended. Since it is desirable to suspend the performance relatively quickly when interference with a foreground

process is detected, m should be small, and accordingly α should be large.

By definition, α is the probability that the normal performance will be judged incorrectly to be degraded.

5 At the extreme value of $\alpha = 1$, the performance is always judged to be degraded, irrespective of any measurement. As described above, the suspension time grows exponentially with each judgement of degradation, and is reset to its initial value when judged normal. In other words, during times of normal performance, the tasks suspension state is a birth-death system that is isomorphic to a bulk-service queue (L. Kleinrock, *Queuing Systems Volume 1: Theory*, John Wiley and Sons, (1975)) of infinite group size with an arrival rate of α and a bulk service rate of β . Therefore, under an assumption that timing measurements exhibit no low-order serial correlation apart from that due to load, the steady-state probability that exactly k judgements of degradation have occurred since the most recent judgement of normalcy is given by:

$$p_k = [\beta / (\beta + \alpha)] [\alpha / (\alpha + \beta)]^k$$

When the next judgement of degradation occurs, which it will after a minimum of m time slices with a probability of α , the suspension time will be 2^k times the base suspension time d , where d is expressed as a multiple of the task's mean time slice. Thus, the mean steady-state inactive time relative to the working time is:

$$q = \sum_{k=0}^{\infty} \alpha 2^k d m^{-k} p_k = (d/m) [\alpha \beta / (\beta - \alpha)]$$

As can be appreciated from the above formula, α needs to be less than β for this system to be stable. Increasing β relative to α increases the duty cycle of the background task during periods without foreground activity. Values for α and β have been selected as $\alpha = 0.05$ and $\beta = 0.2$, whereby $m = 5$ and a 200-millisecond time slice yields a minimum reaction time of one (1) second. A base suspension time of five (5) seconds makes $d = 25$, giving a background performance penalty of $q = 0.33$. Tested empirically, these values provide a relatively prompt reaction to foreground activity, a moderately stable background state, and a fairly low impact on foreground processing in the presence of highly stochastic operation timing.

Turning to another consideration of the present invention, a single process may include more than one background task. However, if multiple tasks use a common resource, then the operation of one task may interfere with the performance of another task, leading the other task to incorrectly suspend its operation. This may lead to unfair measurements, instability, and moreover, two or more background tasks may mutually interfere in a manner that forces each other into their maximally suspended states, even at times when there are no foreground processes in execution.

The background task controller 112 lets only one task execute at a time to overcome the problems with multiple tasks. To determine which task to run, relative priorities are established for the tasks, and each task provides a value that indicates its current priority, referred to as its instantaneous task priority. The background task controller 112 uses the instantaneous task priority to calculate an effective task priority using the following formula:

$$\text{effective task priority} = \text{instantaneous task} \\ \text{priority} \cdot (\text{task duty cycle} + \epsilon)$$

5 where the task duty cycle is calculated as the fraction
of recent real time that the task has been permitted to
execute, exponentially averaged over a period determined
by a settable time constant. The ϵ variable is a small
positive value, ensuring that only a zero instantaneous
10 task priority yields a zero effective task priority. The
task having the smallest value for its effective task
priority from among those tasks that are ready to execute
is the task that is permitted to execute next. Each task
is thus permitted to execute for a portion of real time
15 that is inversely proportional to the value of its
instantaneous task priority.

Task priority is only relevant when two or more
tasks are ready to execute at the same time, and thus
tasks that are throttled or that have naturally low duty
20 cycles do not significantly interfere with other tasks,
even when such tasks have relatively high priorities.
Tasks need to set their instantaneous task priorities so
that they do not execute too frequently and prevent other
tasks from executing.

25 By way of example, each groveling task of the SIS
groveler 60 calculates its priority according to the
following formula:

$$\text{instantaneous task priority} = \text{partition's disk} \\ \text{free space} / \text{total free disk space} + \epsilon$$

30

where ϵ is a small constant.

As a result, each groveling task has a duty cycle
that is inversely proportional to its partition's share
of the total free disk space in the system. At an

extreme, a partition with no free disk space has a groveling task with an instantaneous task priority of ϵ , whereby it has a higher priority than other tasks except those with a zero instantaneous task priority.

5 Note that the above mechanism is directed to background tasks within a process, but may be extended to apply to background tasks within separate processes. Although there is no adequate way for a process to know when it is more important than another process, the
10 process may inform the system when it is less important than usual. A process may set an instantaneous process priority, (where one is normal, zero is minimal), whereby the background task controller calculates an effective process priority using the following formula:

15 effective process priority = instantaneous
 process priority \cdot (1 - process duty cycle)

where the process duty cycle is analogous to the task
20 duty cycle described above, exponentially averaging the fraction of recent real time that the process has been permitted to execute. Of those processes that are ready to execute a task, the one having the highest effective process priority is permitted to execute next, whereby
25 each normal-priority process will be granted an equal share of real time in which to execute. Lower priority processes will receive proportionally reduced shares of execution time.

As can be seen from the foregoing detailed
30 description, there is provided a method and system that regulate a background process based on actual performance measurements. The method and system enable the background process to execute, yet in a manner that limits its interference with a foreground process, even

when the limiting resource that is causing the interference is not the CPU.

While the invention is susceptible to various modifications and alternative constructions, certain
5 illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific form or forms disclosed, but on the contrary, the intention is to cover all
10 modifications, alternative constructions, and equivalents falling within the spirit and scope of the invention.

WHAT IS CLAIMED IS:

1. A computer-readable medium having computer-executable instructions, comprising, executing a background task, receiving data indicative of a measured progress of the background task, and determining when to run the background task based on the data.
2. The computer-readable medium of claim 1 having further computer-executable instructions for suspending the background task for a suspend time, and authorizing the background task to execute.
3. The computer-readable medium having computer-executable instructions of claim 1, wherein determining when to run the background task includes comparing the measured progress of the background task against a target progress.
4. The computer-readable medium having computer-executable instructions of claim 3, wherein determining when to run the background task includes determining a suspend time for suspending the background task.
5. The computer-readable medium having computer-executable instructions of claim 4, wherein if the measured progress with respect to the target progress is acceptable, setting the suspend time to a relatively low duration.
6. The computer-readable medium having computer-executable instructions of claim 5, wherein setting the suspend time to a relatively low duration includes setting the suspend time to a minimum value.

7. The computer-readable medium having computer-executable instructions of claim 4, wherein if the measured progress with respect to the target progress is not acceptable, increasing the suspend time.

5

8. The computer-readable medium having computer-executable instructions of claim 7, wherein increasing the suspend time includes doubling a previous suspend time.

10

9. The computer-readable medium having computer-executable instructions of claim 4 wherein if the measured progress with respect to the target progress is not determinable by present data, maintaining the suspend time.

15

10. The computer-readable medium having computer-executable instructions of claim 1, wherein determining when to run the background task includes determining a suspend time for suspending the background task, and wherein the suspend time is further based on a relative importance of the task.

20

11. The computer-readable medium having computer-executable instructions of claim 1, wherein determining when to run the background task includes statistically combining the data received with previous data.

25

12. The computer-readable medium of claim 1 having further computer-executable instructions for, determining a target value indicative of a target amount of work, and wherein determining when to run the background task includes comparing the measured progress of the background task against the target progress.

30

13. The computer-readable medium having computer-executable instructions of claim 12, wherein determining a target amount further comprises, using measured
5 progress data to automatically calibrate the target amount.

14. The computer-readable medium having computer-executable instructions of claim 1, wherein the measured
10 progress comprises an amount of work performed per unit time.

15. The computer-readable medium having computer-executable instructions of claim 1, wherein the
15 background task is executed for a limited time, and the data include a count of the number of operations performed during the limited time.

16. The computer-readable medium having computer-executable instructions of claim 15, wherein the data
20 include the total real time taken for the operations to complete.

17. The computer-readable medium having computer-executable instructions of claim 16, wherein the data
25 include a parameter representing the relative amount of work performed by each operation.

18. The computer-readable medium having computer-executable instructions of claim 17, wherein the
30 background task is part of a process for recognizing duplicate files on a file system partition, and wherein the amount of work performed by each operation is an amount of data read from the partition.

19. The computer-readable medium having computer-executable instructions of claim 1, wherein the background task performs input/output operations on a resource.

20. The computer-readable medium of claim 1 having further computer-executable instructions for prioritizing an execution of at least one other task.

21. A system for regulating the execution of a background task, comprising, a task regulator for authorizing the background task to execute, a performance evaluator for receiving measured progress data from the background task and providing information corresponding thereto, and a computation mechanism connected to receive the information from the performance evaluator, the task regulator connected to the computation mechanism for authorizing the background task to execute based on information received from the computation mechanism.

22. The system of claim 21 wherein the information received from the computation mechanism includes a suspension time.

23. The system of claim 22 wherein the computation mechanism further determines the suspension time based on a relative importance of the task.

24. The system of claim 21 wherein the performance evaluator compares the measured progress data against target progress data.

25. The system of claim 24 further comprising a calibration mechanism for providing the target progress data.

5 26. The system of claim 25 wherein the calibration mechanism receives the measured progress data and statistically combines the measured progress data with previous progress data.

10 27. The system of claim 21 wherein the measured progress comprises an amount of work performed per unit time.

15 28. The system of claim 21 wherein the background process executes for a limited time, and wherein the measured progress data include a count of the number of operations performed during the limited time and a total real time taken for the operations to complete.

20 29. The system of claim 21 wherein the progress data further include a parameter representing a relative amount of work performed by each operation.

25 30. The system of claim 21 further comprising a resource accessed by the background process.

30 31. The system of claim 21 further comprising at least one other task having its execution regulated by the task regulator.

32. A method of executing a background process, comprising the steps of, executing a task of the background process, measuring the progress of the task, comparing the progress of the task against a target

progress, and if the progress of the task is degraded relative to the target progress, increasing a delay time from a previous value thereof, and suspending for the delay time before re-executing the task.

5

33. The method of claim 32 wherein the progress of the task corresponds to an amount of work per unit time performed by the task.

10

34. The method of claim 32 wherein the background process seeks duplicate files of a file system.

35. The method of claim 32 further comprising the step of adjusting the target amount based on the progress
15 of the task.




FIG. 1




FIG. 2A




FIG. 2B

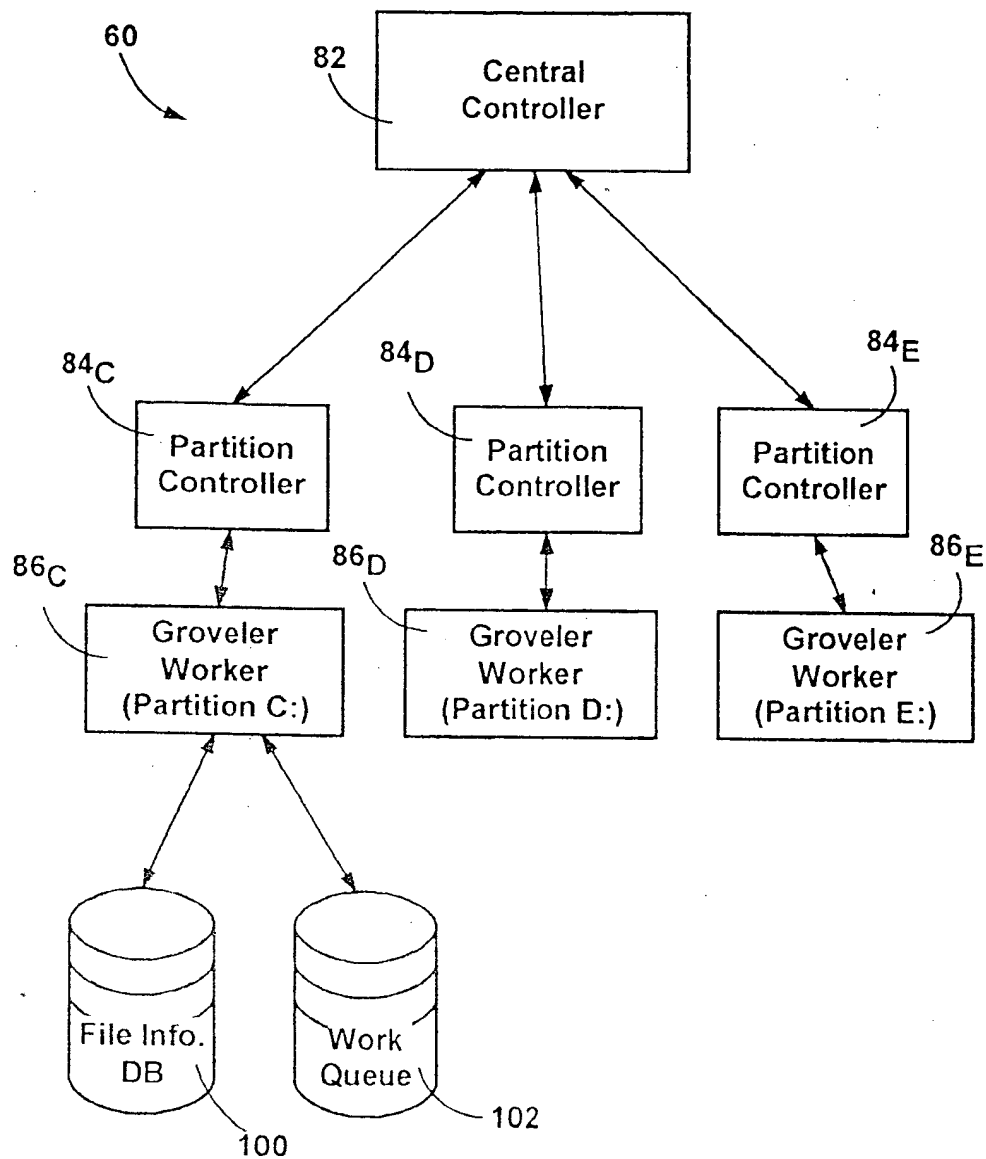



FIG. 3






FIG. 5

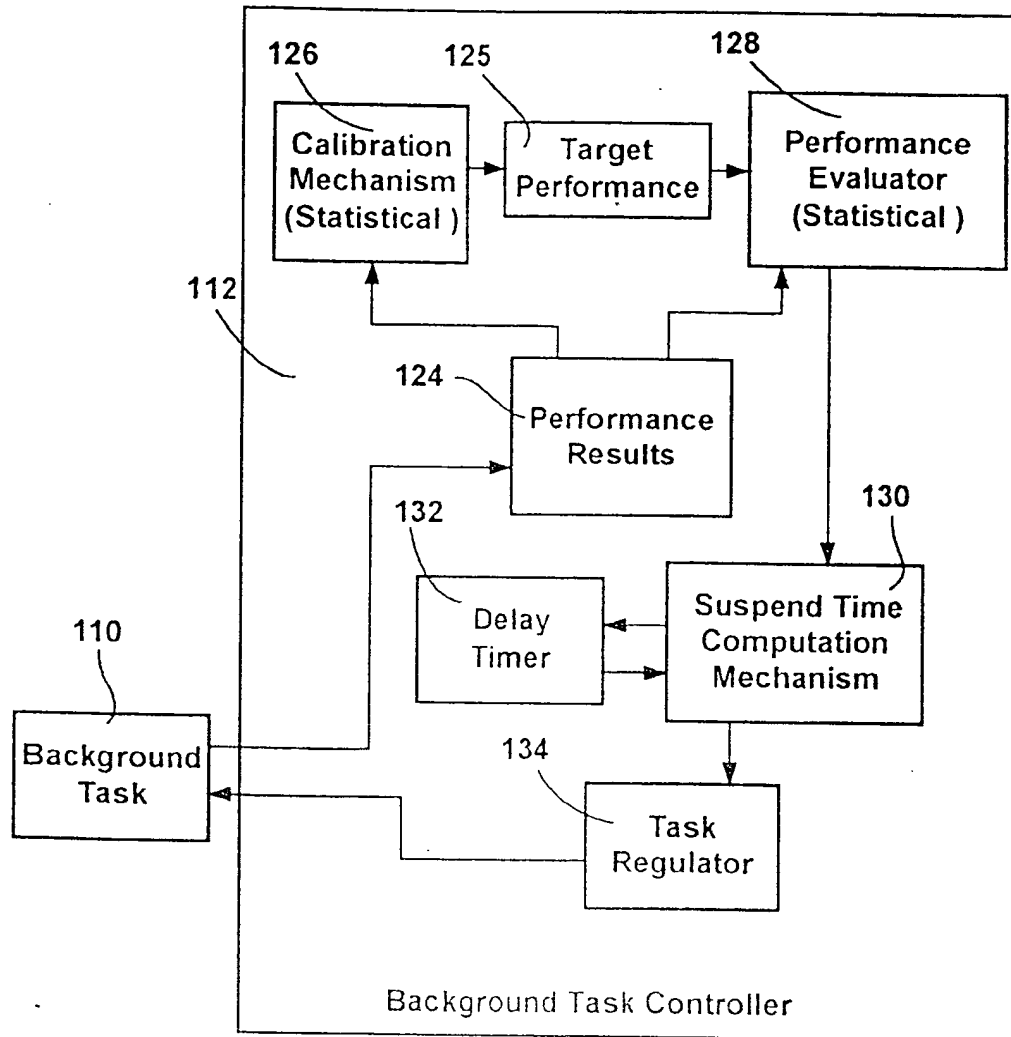
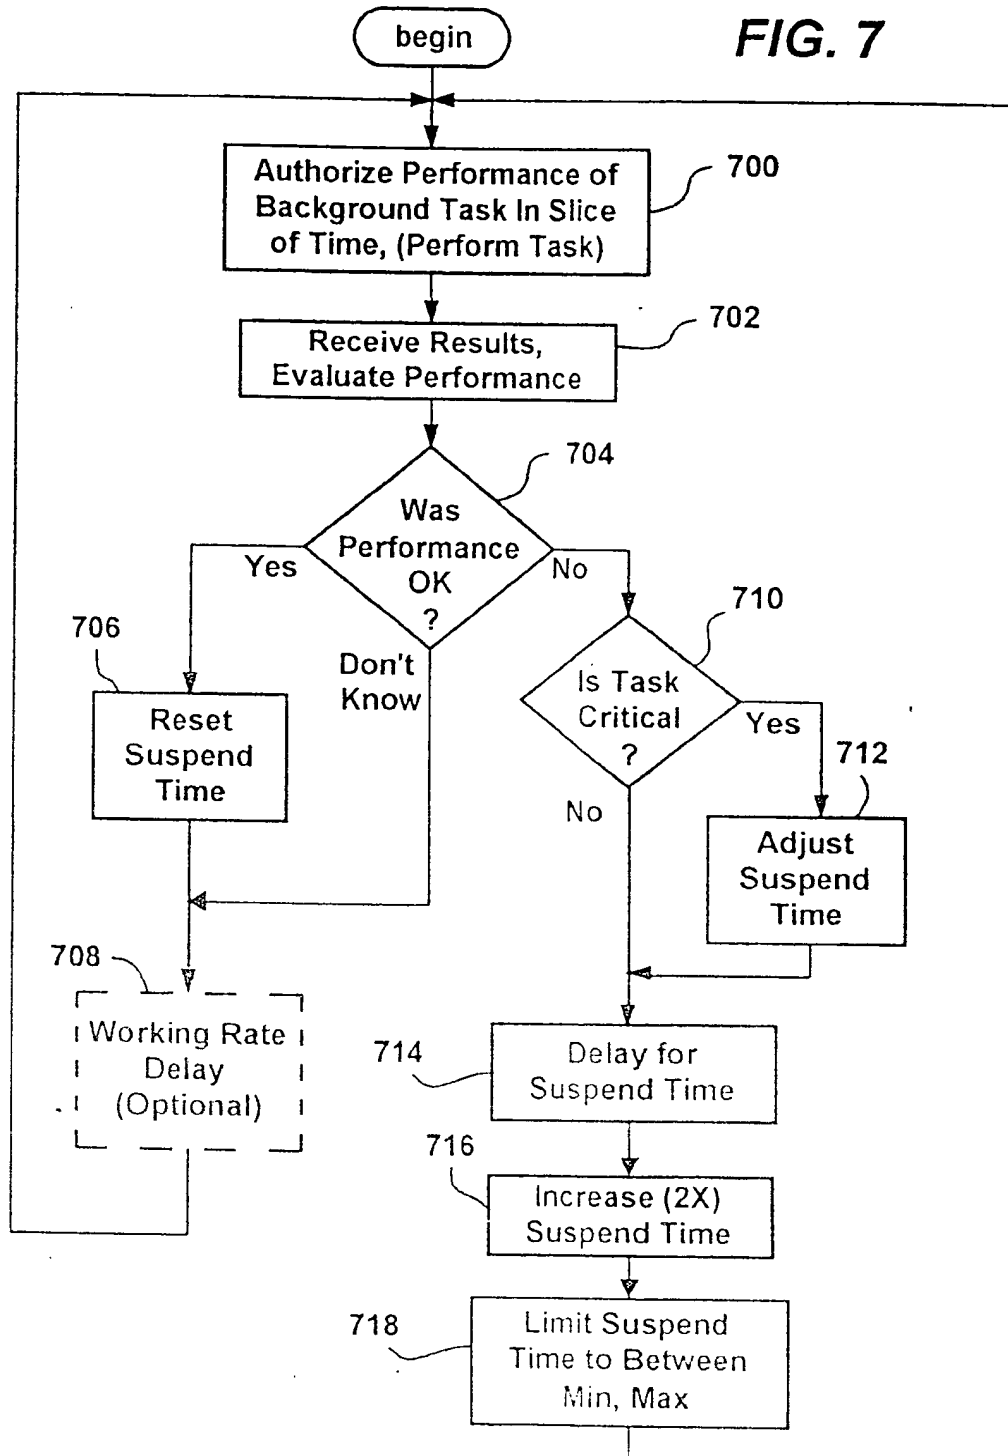


FIG. 6

FIG. 7



INTERNATIONAL SEARCH REPORT

International Application No.

PCT/US 00/18989

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 99 21082 A (IBM) 29 April 1999 (1999-04-29) page 8, line 31 -page 9, line 15 page 27, line 21 -page 28, line 37 page 30, line 6 -page 31, line 30 ---	1, 14, 15, 19-21, 27, 30, 31
P, X	DOUCEUR J R ET AL: "Progress-based regulation of low-importance processes" 17TH ACM SYMPOSIUM ON OPERATING SYSTEMS PRINCIPLES, KIAWAH ISLAND RESORT, SC, USA, 12-15 DEC. 1999, vol. 33, no. 5, pages 247-260, XP002155706 Operating Systems Review, Dec. 1999, ACM, USA ISSN: 0163-5980 the whole document ---	1-35
	-/--	



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

18 December 2000

Date of mailing of the international search report

05/01/2001

Name and mailing address of the ISA

European Patent Office, P.O. Box 18 Patentkan 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2010, Ex. 31 651 ext. nl
Fax: (+31-70) 340-3016

Authorized officer

Carciofi, A

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/18989

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>STEERE D C ET AL: "A feedback-driven proportion allocator for real-rate scheduling"</p> <p>THIRD SYMPOSIUM ON OPERATING SYSTEMS DESIGN AND IMPLEMENTATION, NEW ORLEANS, LA, USA, 22-25 FEB. 1999, vol. spec. issue., pages 145-158, XP002153159</p> <p>Operating Systems Review, Winter 1998, ACM, USA</p> <p>ISSN: 0163-5980</p> <p>abstract; figure 1</p> <p>page 145, right-hand column, paragraph 2</p> <p>-page 146, left-hand column, paragraph 3</p> <p>page 147, left-hand column, paragraph 1 - paragraph 2</p> <p>page 148, left-hand column, paragraph 3</p> <p>-page 151, left-hand column, paragraph 1</p> <p>page 153, right-hand column, paragraph 2</p> <p>-page 154, right-hand column, paragraph 3</p> <p>---</p>	1-35
A	<p>KEN LALONDE: "Batch daemon - README"</p> <p>UNIX BATCH COMMAND (UNIVERSITY OF TORONTO), 'Online!'</p> <p>27 February 1997 (1997-02-27), XP002155707</p> <p>Toronto (CANADA)</p> <p>Retrieved from the Internet:</p> <p><URL:ftp://ftp.cs.toronto.edu/pub/batch.tar.Z> 'retrieved on 2000-12-08!'</p> <p>page 1, line 4 - line 9</p> <p>page 2, line 9 - line 30</p> <p>-----</p>	1-35

...formation on patent family members

PCT/US 00/18989

Form PCT/ISA/210 (patent family annex) (July 1992)